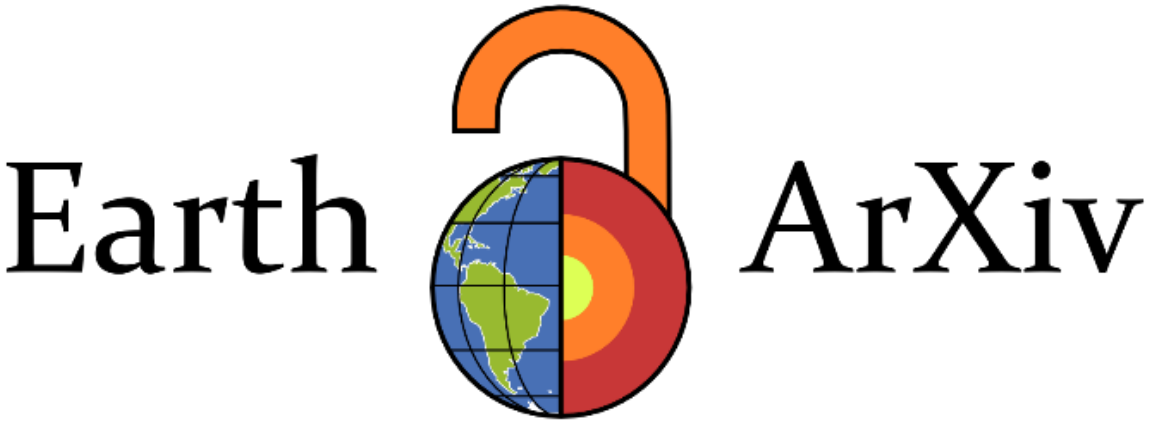

Peer review status:

This is a non-peer-reviewed preprint submitted to EarthArXiv.

Please note that the manuscript has not yet been formally accepted for publication. Subsequent versions may contain slight differences in content. If accepted, the final version will be available via the "Peer-reviewed Publication DOI" link on the right-hand side of this webpage. Please feel free to contact the corresponding author, we welcome your feedback.

Two Wind Farms, Two Islands: Physics-Informed Causal Wind Analysis in New Zealand

Gururaj H C[1*], Dr Vasudha Hegde [2]

[1]Independent Researcher and Data Scientist, Davangere, India
[2]Professor and Head, Department of Electrical and Electronics Engineering, Jain University, Bengaluru, India
[*]Corresponding author: gururaj008@gmail.com

HIGHLIGHTS

- Double Machine Learning corrects the misleading negative correlation between air density and wind power found in raw observational data.
- A 0.1 kg/m³ increase in air density increases power output by about 17.2 MW at the Te Apiti wind farm (North Island) and 3.9 MW at the White Hill wind farm (South Island), implying Te Apiti is about 4.4× more sensitive to density changes.
- The density effect is strongest during summer and at medium wind speeds (6–10 m/s), providing a meaningful operational margin under marginal wind conditions.
- The study uses up to 20 years of hourly ERA5 reanalysis and New Zealand grid wind generation data (2005–2024) to estimate and validate these causal effects across distinct climate regimes.

ABSTRACT

Wind power forecasting models often rely on correlation-based methods, which can misinterpret the relationship between meteorological variables and power generation. A key example is air density: while physics suggests denser air should increase available wind power, observational data can show a negative correlation because high-pressure regimes (high density) often coincide with lower wind speeds. This study applies a Double Machine Learning (DML) framework to separate these confounding effects using up to 20 years of hourly ERA5 reanalysis and New Zealand grid wind generation data (2005–2024). Because it is not feasible to study every wind farm, we use one representative wind farm per island—Te Apiti for the North Island and White Hill for the South Island—so the estimates should be interpreted as site-based evidence for island-specific operating regimes rather than island-wide fleet averages. Using XGBoost nuisance models with time-ordered cross-fitting, and controlling for wind dynamics and diurnal/seasonal structure via engineered covariates (including wind speed, lag/ramp terms, and harmonic seasonality features), we mitigate the spurious negative association seen in raw data. The results show a positive causal effect: a 0.1 kg/m³ increase in air density raises power output by about 17.2 MW at the Te Apiti wind farm (North Island) and 3.9 MW at the White Hill wind farm (South Island). The larger North Island effect is consistent

with regional differences in operating conditions and fleet characteristics, but attribution to specific turbine control technology should be treated as a hypothesis. Subgroup estimates indicate the density effect is strongest in summer and at medium wind speeds (6–10 m/s), which is relevant for operational planning under marginal wind conditions.

KEYWORDS

Causal inference; Double machine learning; Wind generation; Air density; ERA5 reanalysis; Wind power forecasting; Heterogeneous treatment effects; Confounding; SHAP; New Zealand

1. Introduction

The global transition toward decarbonized energy systems has accelerated the deployment of wind power as a primary source of renewable electricity. As wind penetration increases, the inherent variability of wind resources creates challenges for grid stability and market operations. This makes accurate modelling and forecasting of wind generation from meteorological conditions a basic requirement for secure power system operation. However, high forecast accuracy alone does not guarantee that a model has learned the true physical drivers of generation, which motivates a causal (not purely correlational) weather–power attribution framework.

1.1. Motivation: why causal (not correlational) weather–wind-power attribution

New Zealand's electricity system is increasing renewable penetration, and this makes short-term management of intermittent resources like wind a day-to-day operational issue. The grid relies on wind generation forecasts to balance supply and demand, and the Electricity Authority of New Zealand has highlighted the need for improved wind generation modelling to reduce reliance on expensive thermal backup [1]. Many current forecasting pipelines use machine learning models that prioritise predictive accuracy, but they provide limited physical interpretability and can behave poorly under unusual weather regimes.

A core limitation is that these approaches often learn correlation rather than causation. Meteorological variables are strongly correlated with each other, so a model can fit relationships that are statistically strong but physically misleading. In this study, air density illustrates the issue: physics suggests higher air density should increase power output, but raw observational data can show a negative density–power correlation because high-pressure (high-density) conditions can coincide with calmer winds. Pearl's causal framework highlights why this happens: without an explicit structural model of the data-generating process, associations can be misread as drivers, leading to incorrect inferences and operational decisions [2].

This motivates a causal inference approach rather than another predictive regression. Using Directed Acyclic Graphs (DAGs), the analysis can make assumptions explicit and identify confounders (for example, seasonality and synoptic-scale patterns) that must be adjusted to estimate the marginal effect of air density on generation [3]. This is also relevant for long-term planning, because correlations can change as climate patterns shift, whereas the causal mechanisms tied to turbine physics are expected to remain stable.

1.2. Research gap and contributions

The basic physics of wind power implies that available power increases with air density and roughly with the cube of wind speed, within fundamental conversion limits [4]. Despite this, most operational forecasting work prioritises wind-speed prediction, and other atmospheric variables (for example temperature and pressure) are often treated in a correlational way rather than as causal drivers [5]. This becomes a problem because meteorological variables are confounded: synoptic regimes that raise pressure (and air density) can also reduce wind speeds, so purely data-driven models can learn relationships that look strong in-sample but do not reflect the underlying mechanism.

A second gap is that many studies treat the wind fleet as a single homogeneous system, even though turbine control and operating regimes can change how meteorological variability translates into power. For example, variable-speed turbines with pitch control can regulate output differently from older designs, which can alter sensitivity to air-density-related effects at the same wind speed [6]. If these differences are ignored, model coefficients and feature importance can mix multiple behaviours and become hard to interpret for system-level planning.

This study addresses these gaps through three contributions. First, we apply a Double Machine Learning framework to estimate the causal effect of air density on generation while adjusting for weather-driven confounding. Second, we analyse the Te Apiti wind farm (North Island) and the White Hill wind farm (South Island)wind generation separately to quantify regional differences in the estimated density effect and to test for heterogeneity consistent with differences in fleet and operating conditions. Third, we report where the effect is largest across practical operating regimes (season and wind-speed bands), so the results can inform when air density is most relevant for grid operations beyond simple correlation.

1.3. Paper organization

The remainder of this paper is structured as follows. Section 2 describes the data acquisition and preprocessing pipeline, including the integration of ERA5 reanalysis with grid generation data and the derivation of physics-informed features. Section 3 presents exploratory diagnostics, highlighting non-stationarity in wind generation and confounding patterns in raw

correlations. Section 4 formalizes the causal framework by defining the Directed Acyclic Graph (DAG) and the identification strategy used to isolate the effect of air density. Section 5 describes the Double Machine Learning (DML) methodology, including nuisance model specification and cross-fitting. Sections 6 and 7 present the main causal estimates, heterogeneity analyses, and validation using robustness and sensitivity checks. Section 8 discusses operational implications for grid stability, and Section 9 concludes with limitations and future work.

2. Data and Study Design

To estimate the causal drivers of wind power, we constructed a longitudinal dataset spanning 20 years (2005–2024). This study integrates high-frequency wind generation logs from the New Zealand Electricity Authority with meteorological reanalysis from ERA5. The following sections describe the study area and the Te Apiti wind farm (North Island) vs the White Hill wind farm (South Island)context, the quality control and unit correction applied to the generation data, and the physics-informed feature construction used to create the final analysis dataset.

2.1 Study Area: The Te Apiti wind farm (North Island) vs White Hill wind farm (South Island) contexts

New Zealand sits in the mid-latitude westerly flow, but local terrain and exposure create different wind-generation regimes across the two islands [8][9]. This study compares a representative North Island wind site (Te Apiti) and a representative South Island wind site (White Hill) to test whether the estimated weather–power relationships, especially for air density, behave similarly under different geographic settings.

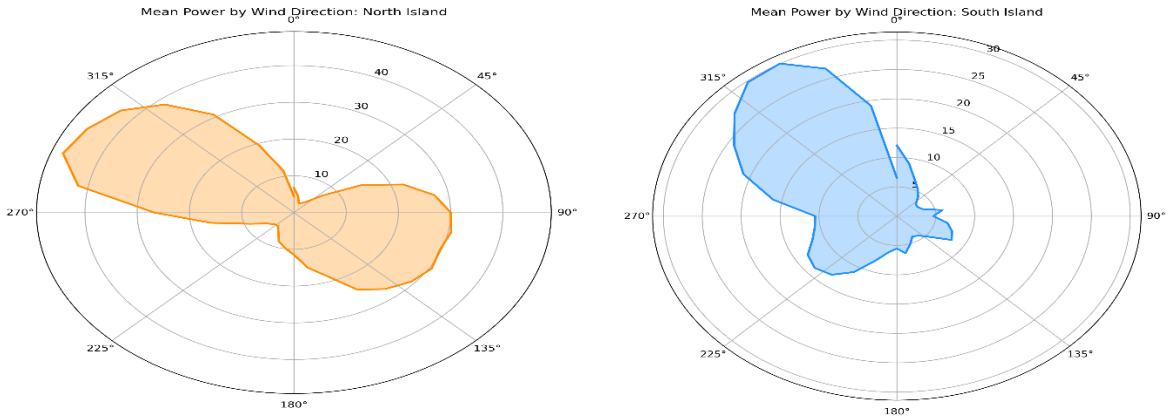

Figure 1. *Directional dependence of wind power generation across New Zealand's distinct aerodynamic regimes.* ***(A)*** *North Island (Te Apiti) exhibits a bi-directional channelling effect (bow-tie pattern) driven by the local topography of the Manawatū*

Gorge. ***(B)*** *South Island (White Hill) displays a dominant uni-directional westerly flow characteristic of the "Roaring Forties" planetary circulation. Radial axes represent mean power output (MW).*

For the Te Apiti wind farm (North Island), the mean-power polar plot (**Figure 1A**) shows a strongly directional pattern with two dominant lobes: one in the westerly sector around 270°–330° and another in the easterly sector around 90°–140°. This indicates that high generation is concentrated in a narrow set of flow directions, consistent with terrain channelling effects that can occur in New Zealand's complex topography [8].

For the White Hill wind farm (South Island), the corresponding polar plot (**Figure 1B**) shows a broader dominant lobe from the west to south-west sector (roughly 270°–330°), with weaker generation in most other directions. This pattern is consistent with stronger exposure to large-scale westerly circulation in the lower South Island, where synoptic variability can dominate over local channel effects [9].

Comparing these two contexts is important for the causal analysis because it checks whether the air-density effect is stable across a terrain-influenced regime (Te Apiti wind farm in North Island) and a more open, westerly-exposed regime (White Hill wind farm in South Island).

2.2 Wind Generation Data: Resolution, QC, and Unit Correction

Wind generation data were sourced from the New Zealand Electricity Authority (EMI) for January 2005 to December 2024 [1]. The raw series is provided at 30-minute trading-period resolution as energy (kWh per trading period). To match the hourly meteorological inputs, we aggregated two consecutive trading periods into one hour by summing their energy and converting to hourly average power. This yields power in kW because 1 hour of energy in kWh is numerically equal to average kW over that hour.

Table 1: *Data summary and quality checks for the North Island (Te Apiti) and South Island (White Hill) hourly wind generation datasets (observation period, sample size, physical ranges, missingness, and wind–power correlation)*

Metric	North Island (Te Apiti)	South Island (White Hill)
Observation Period	Jan 2005 – Dec 2024	Jan 2009 – Dec 2024
Total Observations (Hours)	175308	140256
Max Power Output	88.96 MW	55.37 MW
Min Power Output	0.00 MW	0.00 MW
Max Wind Speed (100m)	22.65 m/s	30.40 m/s
Temperature Range	-4.9°C to 29.4°C	-15.1°C to 26.8°C
Missing Values	0 (0.0%)	0 (0.0%)
Wind-Power Correlation	0.7124	0.6608

A key quality-control step was unit validation. Before correction, some values exceeded 30,000, which is not plausible for wind farms with installed capacities of roughly 90 MW (Te Apiti) and 58 MW (White Hill). After confirming the raw values were in kW, we divided by 1,000 to convert to MW. The post-correction maximum observed outputs were 88.96 MW

(North Island) and 55.37 MW (South Island), consistent with the expected upper bounds reported in the QA summaries (**Table 1**).

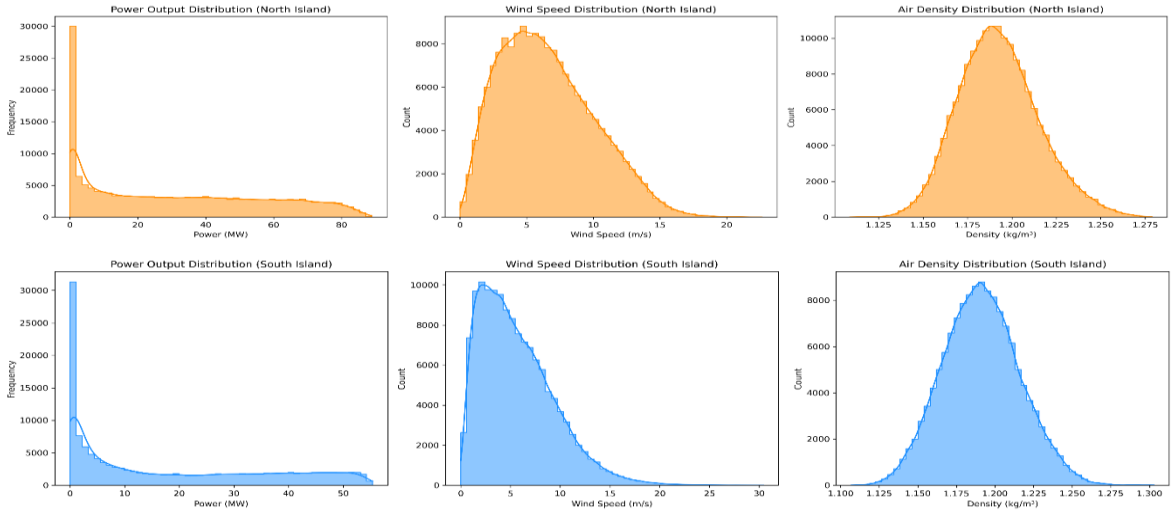

Figure 2. *Univariate distributions for the main variables used in the causal analysis. The top row shows North Island (Te Apiti) hourly distributions of (left) wind power output (MW), (middle) 100 m wind speed (m/s), and (right) air density (kg/m³). The bottom row shows the same three distributions for the South Island (White Hill). These plots summarise the marginal variability and highlight the strong mass near zero generation alongside the long right tails in wind speed and power that motivate careful confounding control in later causal modelling.*

Figure 2 shows the resulting power distributions for both sites. The representative sites from both the islands have a strong mass near 0 MW (calm hours and/or curtailed hours), with a long upper tail up to each site's maximum output. Negative values, which are non-physical for generation, were removed by clipping to zero so that the final series respects the non-negativity constraint used in later causal modelling.

2.3 Meteorological Data and Feature Construction

Meteorological predictors were obtained from the ERA5 reanalysis produced by the European Centre for Medium-Range Weather Forecasts (ECMWF)[10]. For each wind farm (Te Apiti in the North Island and White Hill in the South Island), we extracted hourly ERA5 fields at the farm location and aligned them to the generation time index. Wind conditions were represented using the 100 m zonal and meridional wind components (u and v), which were converted to hub-height wind speed and wind direction. We also used near-surface temperature and surface pressure to represent the thermodynamic state relevant for air-density estimation.

Beyond raw variables, we constructed physics-informed features that map more directly to wind energy conversion. The key derived variable is air density (ρ), computed from surface pressure and temperature using the ideal gas relationship $\rho = P/(R \cdot T)$. This matters because changes in air density shift the kinetic energy flux through the rotor even when wind

speed is unchanged, but density is often treated as constant in operational models [11]. By explicitly including ρ, the causal model can separate “how much air is moving” (wind) from “how heavy the air is” (thermodynamics).

Figure 2 summarises the marginal distributions of wind speed and air density for the representative sites from both islands. Wind speeds show the expected right-skewed shape, while air density is approximately bell-shaped, with the White Hill wind farm (South Island) distribution slightly shifted toward higher density relative to the Te Apiti wind farm (North Island). This is consistent with the broader point that density varies seasonally and regionally and should be considered when estimating wind energy potential and capacity factors [11].

To capture key non-linearities and operational dynamics, we engineered additional predictors used in the main analysis: windcubed (wind speed cubed), turbulenceproxy (rolling standard deviation of wind speed over a short window), ramprate (hour-to-hour change in wind speed), and windlag1 (one-hour lag of wind speed). These features help the model distinguish steady productive winds from volatile regimes associated with ramping and potential cut-out behaviour, improving realism of the estimated effects under grid-relevant operating conditions.

2.4 Final Analysis Dataset

The final analysis dataset was created by merging the hourly ERA5 features with the processed hourly generation series on the timestamp index. A key preprocessing step was time alignment: ERA5 timestamps were converted from UTC to New Zealand local time so that the meteorological drivers line up with the generation logs. As a sanity check, the diurnal temperature cycle was inspected to confirm that daily peaks occur in local afternoon hours after conversion.

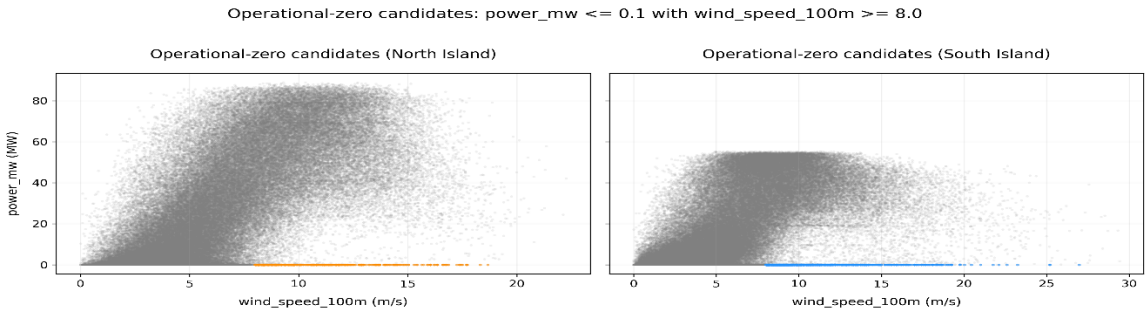

Figure 3. *Operational-zero candidates (power_mw ≤ 0.1 with wind_speed_100m ≥ 8.0) highlighted on the wind speed–power scatter for North and South Island.*

A practical challenge is separating low output due to low wind (physics) from low output due to curtailment, outages, or operational constraints. **Figure 3** highlights a clear set of

"operational-zero candidates", defined as hours where power is near zero (power_mw ≤ 0.1) despite strong winds (wind_speed_100m ≥ 8.0 m/s). In a causal setting, keeping many such points can create limited overlap (positivity issues) for the active-generation regime, because the outcome is effectively fixed near zero over a range of high-wind conditions for non-physical reasons [12].

Table 2: *Example hours flagged as operational-zero candidates, showing high wind speed with near-zero power output (NZST timestamps).*

Timestamp (NZST)	Region	Wind Speed (m/s)	Power Output (MW)	Air Density (kg /m ³)
28-12-2010 09:00	Te Apiti wind farm (North Island)	19.08	0	1.13
23-09-2010 03:00	Te Apiti wind farm (North Island)	18.67	0	1.19
28-12-2010 08:00	Te Apiti wind farm (North Island)	18.35	0	1.13
20-09-2023 21:00	White Hill wind farm (South Island)	26.97	0	1.15
05-08-2022 16:00	White Hill wind farm (South Island)	25.22	0.05	1.16
09-12-2023 13:00	White Hill wind farm (South Island)	25.18	0	1.15

To reduce this problem, we excluded operational-zero candidates using the transparent rule above before fitting the main causal models. **Table 2** provides examples of these events, including cases with very high wind speeds (e.g., 26.97 m/s) paired with zero or near-zero power, which is consistent with non-standard operation such as disconnection, maintenance, or high-wind cut-out. This filtering step is intended to focus the target effect on the turbine's physical response during active operation rather than grid-constrained output.

3. Exploratory and Time-Series Diagnostics

Before applying the causal inference framework, we first check the key statistical and physical patterns in the data. This section summarises basic distribution and seasonality diagnostics, then examines bivariate relationships, including the negative air-density–power association observed in the raw data. We also assess time-series behaviour (trend/seasonality, stationarity, and autocorrelation) and review lag structure to guide the lagged and seasonal controls used later in the nuisance models.

3.1 Univariate Distributions & Seasonality

The marginal distributions of the key variables are broadly consistent with expected wind-power behaviour. **Figure 2** shows right-skewed wind speed distributions at both sites, with the White Hill wind farm (South Island) exhibiting a longer high-wind tail than the Te Apiti wind farm (North Island). Power output (**Figure 2**, left panels) is strongly bimodal, with a large mass near 0 MW and a wide spread up to the site maxima, indicating that the dataset includes both low-generation hours and near-capacity operation.

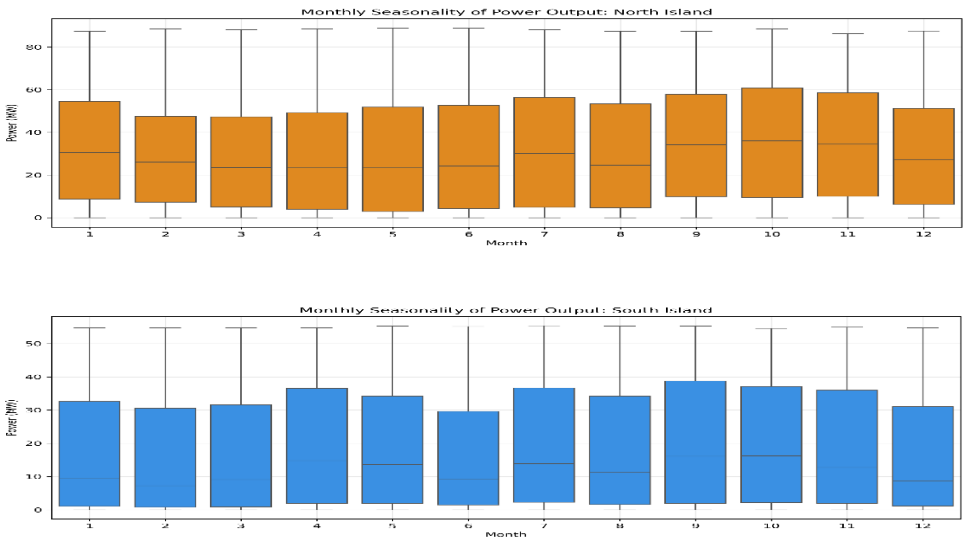

Figure 4. *Monthly seasonality of hourly wind power output for the Te Apiti wind farm (North Island) (top) and White Hill wind farm (South Island) (bottom), shown as boxplots by month.*

As depicted in **Figure 04** both sites show a clear annual cycle, with lower typical generation in late summer (Feb–Mar) and higher typical generation in spring (Sep–Nov). Spring (Sep–Nov) also shows larger within-month spread than late summer (Feb–Mar), consistent with more variable wind conditions during the windier season [9].

North Island (orange/Te Apiti): Medians are generally higher across most months, and the interquartile ranges are typically wider, indicating more sustained production and variability through the year. The spring uplift (Sep–Nov) appears strong and relatively sustained compared with late summer (Feb–Mar).

South Island (blue/White Hill): Typical output is lower (medians closer to zero in several months), while occasional high-output events still occur (long upper whiskers), especially around spring (Sep–Nov). Late summer (Feb–Mar) is among the weakest months.

3.2 Bivariate Physics: Power Curves and the "Negative Density Correlation" Paradox

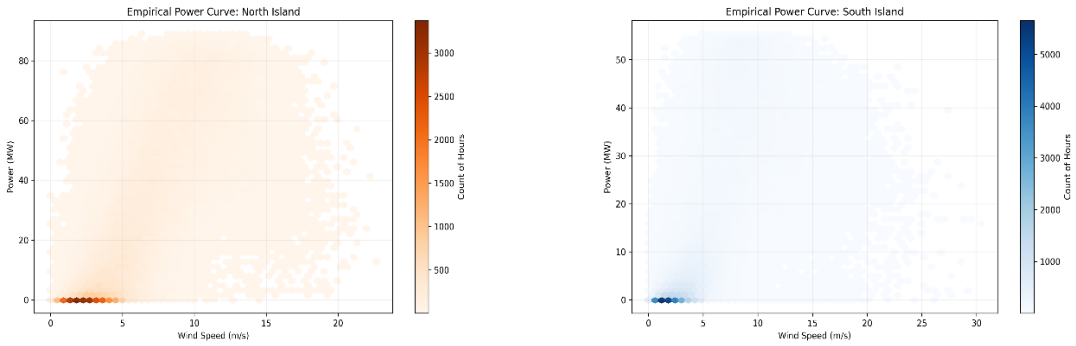

Figure 5. *Empirical wind speed–power curves (Hexbin density) for the Te Apiti wind farm (North Island) (A) and White Hill wind farm (South Island) (B), showing the typical rise in output with wind speed and saturation near rated power.*

The wind speed–power relationship provides a basic physical check for the dataset. **Figure 5A** and **Figure 5B** show the empirical power curves for the Te Apiti wind farm (North Island) and the White Hill wind farm (South Island), where power rises with wind speed in the partial-load region and then levels off near each site's maximum output. This behaviour is consistent with standard turbine power-curve descriptions in wind energy fundamentals [4].

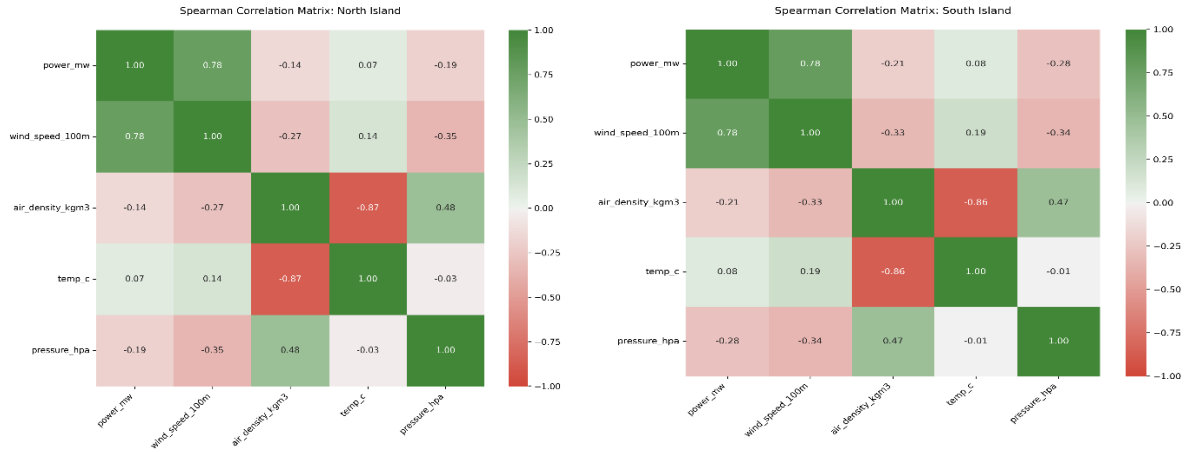

Figure 6. *Spearman correlation matrices for the Te Apiti wind farm (North Island) (A) and White Hill wind farm (South Island) (B) showing the pairwise associations between power output, wind speed (100 m), air density, temperature, and pressure.*

The correlation patterns, however, show a counter-intuitive result for air density. From the Spearman correlation matrices (**Figure 6A** and **Figure 6B**), air density is negatively associated with power (about -0.14 in the Te Apiti wind farm (North Island) and -0.21 in the White Hill wind farm (South Island)), even though basic physics suggests that higher density increases available wind power (proportional to ρv^3). Air density is also negatively correlated with wind speed (about -0.27 North and -0.33 South), meaning density and wind speed tend to move in opposite directions in the raw data.

This pattern is consistent with a confounding/endogeneity problem in a simple bivariate view, where density is entangled with larger-scale weather states that also control wind speed. In that setting, the raw density–power correlation can turn negative even if the underlying physical effect of density on power is positive after conditioning on wind speed and related drivers. This is a standard omitted-variable issue in observational settings and motivates an explicit causal adjustment strategy rather than relying on correlations alone [13].

The combined physics feature theoretical_energy shows a marginally higher Spearman association with power than wind_speed_100m alone in both islands (Te Apiti wind farm (North Island): 0.7215 vs 0.7206; White Hill wind farm (South Island): 0.7189 vs 0.7182). While the numerical difference is small, the direction is consistent across both sites, supporting the interpretation that air density contributes incremental signal once integrated with wind speed, even though its standalone correlation is negative.

3.3 Stationarity & Autocorrelation

Table 3: ADF and KPSS stationarity test p-values for power output and wind speed time series

Region	Variable	ADF p-value	KPSS p-value
Te Apiti wind farm (North Island)	Power	<0.001	0.01
Te Apiti wind farm (North Island)	Wind speed	<0.001	0.01
White Hill wind farm (South Island)	Power	<0.001	0.01
White Hill wind farm (South Island)	Wind speed	<0.001	0.01

To assess time-series assumptions, we ran Augmented Dickey–Fuller (ADF) and KPSS tests on the hourly power and wind speed series for representative cites from both islands (ADF: statsmodels defaults; KPSS: regression='c', nlags='auto'). **Table 3** shows a consistent pattern: ADF strongly rejects a unit root, while KPSS rejects stationarity, indicating that the series are not strictly stationary in the KPSS sense.

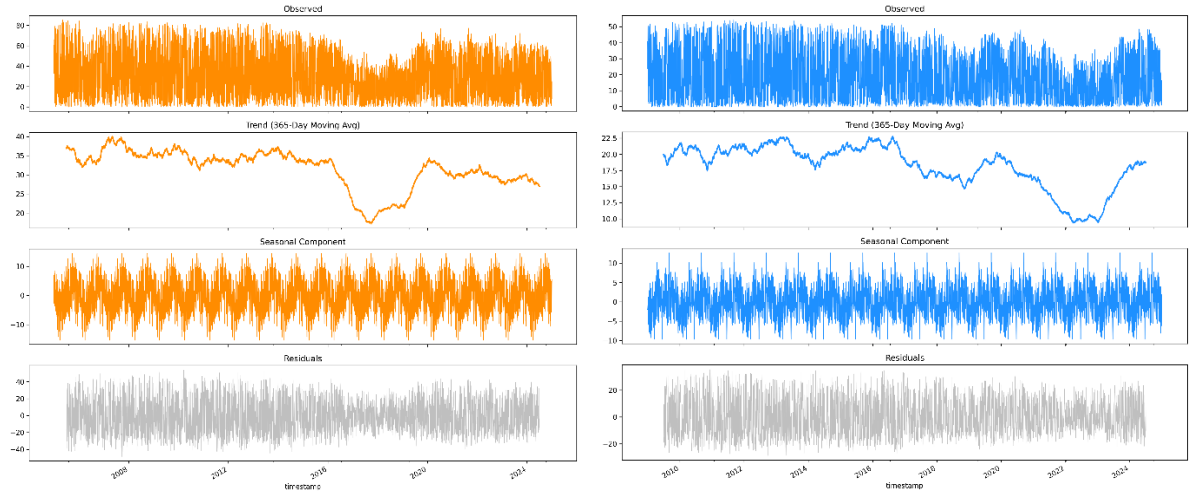

Figure 7: *STL decomposition of daily mean wind power for the Te Apiti wind farm (North Island) (A) and White Hill wind farm (South Island), separating the observed series into a smooth 365-day trend, seasonal component, and residual variability*

This mixed result is consistent with bounded series that contain strong seasonality and slow changes in the mean level over multi-year periods. **Figure 7** (decomposition) shows clear recurring seasonal structure and a smooth low-frequency trend component for both islands, rather than behaviour consistent with an unbounded random walk. Because power output is physically bounded (0 to rated output) and the main non-stationarity appears to be deterministic structure (seasonality + slow trend), we avoid differencing as a default preprocessing step.

Instead, we handle temporal dependence by adding explicit seasonal controls and lagged predictors to the nuisance-model feature set (e.g., sine/cosine harmonics for month/hour and wind lags), so short-term persistence and seasonality are absorbed in the first-stage fits rather

than left in the errors. For inference, we quantify uncertainty using a time-series moving-block bootstrap on the residual-on-residual regression, which is designed to be robust to autocorrelation and heteroskedasticity and plays a similar role to HAC approaches discussed in [14].

3.4 Lag Analysis

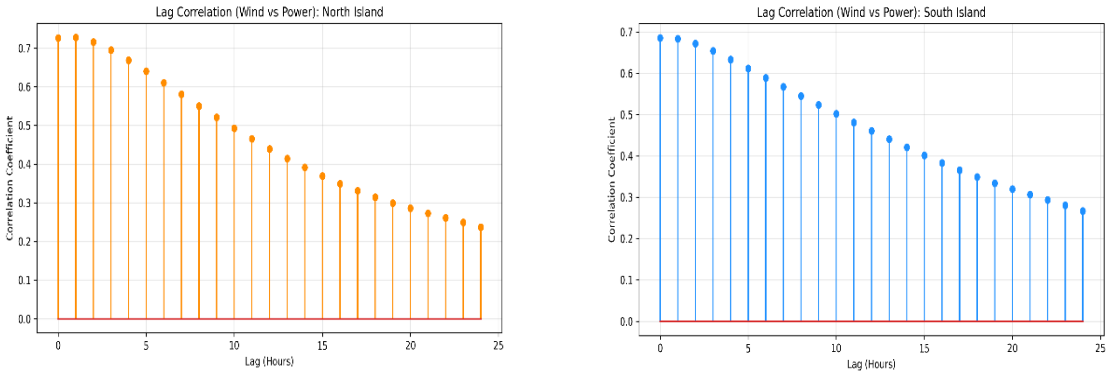

Figure 8. *Cross-correlation between wind speed and power output across hourly lags for the Te Apiti wind farm (North Island) (A) and White Hill wind farm (South Island) (B), showing strong short-lag dependence and a gradual decay over the next 24 hours*

To select a reasonable temporal structure, we computed the cross-correlation function (CCF) between wind speed and power output. **Figure 8** shows the lag correlations for the Te Apiti wind farm in North Island (orange) and White Hill wind farm in South Island (blue). In both regions, the cross-correlation is highest at very short lags (0–1hour) and then decays gradually with increasing lag. The correlation remains materially positive for several hours, indicating persistence in the wind–generation relationship driven by slowly evolving weather patterns.

To account for this short-term memory and reduce the risk of attributing delayed wind effects to other covariates, we include the first lag of wind speed (wind_lag_1) as an explicit predictor in the adjustment set / nuisance feature set. This choice is consistent with standard time-series practice where dependence is handled by including lagged terms and using resampling methods designed for serial correlation [15].

4. Causal Framework

To move from prediction to causal inference, we first state our structural assumptions for wind generation using a Directed Acyclic Graph (DAG). We then use the DAG to define an identification strategy (backdoor adjustment) that isolates the causal effect of air density on power from confounding meteorological structure, including latent synoptic conditions.

4.1 Causal DAG and identification strategy

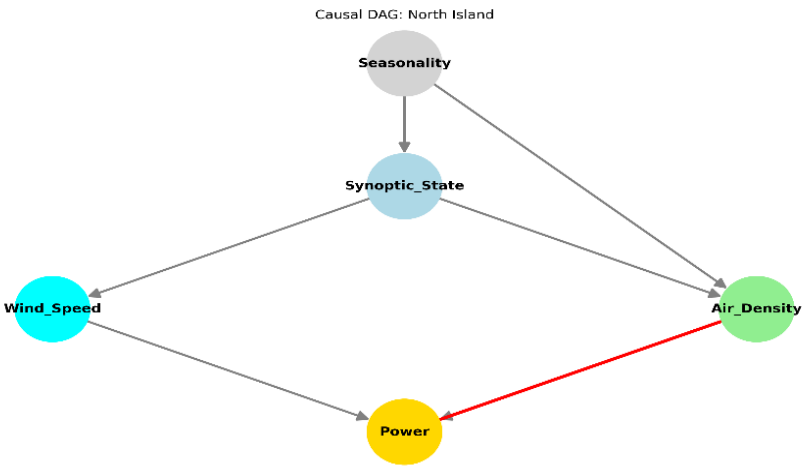

Figure 9: *Causal DAG for the Te Apiti wind farm (North Island) showing the assumed relationships between seasonality, synoptic state, wind speed, air density (treatment), and power output (outcome), with the Air Density → Power link highlighted. The same DAG structure is applied to the White Hill wind farm (South Island), with site-specific parameter estimates obtained separately.*

To identify the causal effect of air density on power generation, we make our structural assumptions explicit using the Directed Acyclic Graph (DAG) in **Figure 9**. The DAG follows the causal-graph framework of Pearl and encodes the physical information flow from large-scale weather drivers to wind speed, air density, and turbine output [2]. Its structure was built from domain knowledge rather than automated discovery, which is recommended in complex observational settings to ensure interpretability and physical plausibility [16].

The main identification challenge is confounding by Synoptic_State, which represents the large-scale weather regime (e.g., anticyclones vs. troughs/fronts). In the DAG, Synoptic_State affects both Air_Density and Wind_Speed, and Wind_Speed is a direct driver of Power. This induces the non-causal association Air_Density ← Synoptic_State → Wind_Speed → Power, which can make high-density conditions appear "bad" for generation simply because they co-occur with calmer winds (as reflected in the negative density–wind-speed association in the exploratory correlations) [17]. Seasonality is included as an upstream driver that shifts the distribution of synoptic regimes and also has a direct pathway to Air_Density.

Our identification strategy uses the Backdoor Criterion to block this open backdoor path. Specifically, conditioning on Wind_Speed closes the path through Synoptic_State, and conditioning on Seasonality removes seasonal structure that influences Air_Density and the occurrence of synoptic regimes [18]. Under these assumptions, the remaining variation in Air_Density after adjustment can be interpreted as contributing to the causal effect on Power, rather than reflecting synoptic confounding.

4.2 Adjustment set derivation

Using the backdoor criterion on the DAG, a valid adjustment set for identifying the Air_Density → Power effect is {Synoptic_State, Seasonality, Wind_Speed}. Since Synoptic_State is not directly observed, we proxy it using cyclical time features (month_sin, month_cos, hour_sin, hour_cos) together with wind_speed_100m, which capture recurring diurnal and seasonal structure in weather regimes that jointly influence density and generation.

We intentionally exclude temperature and surface pressure from the control set because air density is computed from these variables via the ideal gas relationship, so including them would make the treatment almost mechanically determined by the covariates and risks over-control / unstable identification in the causal stage. Instead, we adjust for upstream structure (seasonality and synoptic patterns) and wind speed, which blocks the main confounding paths while keeping usable variation in air density for estimation [3][18]. Finally, given the mixed ADF/KPSS stationarity signals in diagnostics, these time features also help absorb slow-moving temporal structure before applying robust inference in the next phase.

5. Double Machine Learning Methodology

To translate the identification strategy into an estimate, we use the Double Machine Learning (DML) framework. DML uses flexible learners to model the nuisance relationships between weather variables and generation while still enabling valid inference on the target causal effect. The next sections describe the DML workflow, the chosen learners, and the resampling approach used for uncertainty under time dependence.

5.1 DML Overview

We estimate the causal effect of air density on wind power using the Partially Linear Regression (PLR) framework. The model assumes the outcome Y (power generation) is linearly related to the treatment T (air density), while the confounders X (wind speed, seasonality, turbulence, and related controls) can affect both Y and T in a non-linear way:

$$Y = \theta T + g(X) + \epsilon$$

Here, θ is the target causal parameter and $g(X)$ is an unknown function capturing the combined effect of meteorological confounders. In high-dimensional settings, directly estimating θ with regularized or flexible models can lead to biased estimates due to regularization bias [7].

Double Machine Learning addresses this using a residual-on-residual (orthogonalization) approach, also known as Robinson's transformation. First, we estimate the nuisance functions $m(X) = E[Y \mid X]$ and $p(X) = E[T \mid X]$ using machine learning, with cross-fitting to limit overfitting. We then compute residuals $\tilde{Y} = Y - \hat{m}(X)$ and $\tilde{T} = T - \hat{p}(X)$, and finally

regress $\tilde{Y}$ on $\tilde{T}$ using ordinary least squares. This procedure "partials out" the confounding captured in X, enabling valid estimation and inference for θ even when $g(X)$ is complex [19].

5.2 Nuisance Models (XGBoost)

To estimate the nuisance functions $E[Y \mid X]$ and $E[T \mid X]$, we use Extreme Gradient Boosting (XGBoost) as the base learner. XGBoost is well suited for wind power data because it can represent sharp non-linear responses and interaction effects that are hard to capture with linear models, including the saturation behaviour at high wind speeds [20]. We implement XGBoost through the scikit-learn compatible interface to keep the workflow consistent across tuning and cross-fitting [21]. We also rely on prior evidence that gradient boosting performs well for complex prediction tasks in energy and engineering settings [22].

To avoid weak nuisance fits (which can bias DML) while limiting overfitting, we tune key hyperparameters using RandomizedSearchCV. The tuning focuses on parameters controlling model complexity and generalization (e.g., n_estimators, max_depth, learning_rate, subsample, and colsample_bytree), and selects the configuration that minimizes prediction error for the nuisance targets [23].

Because the data are time ordered, we avoid random shuffling when building folds. We use time-aware validation during tuning and non-shuffled splits during cross-fitting to limit look-ahead bias and keep the residuals aligned with the temporal structure of the data.

5.3 Uncertainty Quantification

Standard statistical inference assumes that errors are independent and identically distributed (i.i.d.). In wind generation, this is not realistic because both weather and power output show persistence, so the error at one hour is related to nearby hours. From the ADF and KPSS results in Section 3.3, the series show strong deterministic structure (seasonality and slow trend) along with serial dependence, rather than behaving like an i.i.d. sequence. If this dependence is ignored, standard errors become too small, which makes p-values and confidence intervals too optimistic.

To assess uncertainty under serial dependence, we compute the residual-on-residual estimate θ and its baseline standard error from the final residual regression, and we additionally use a Moving-Block Bootstrap (MBB) on the residual regression as a robustness check [24]. MBB resamples contiguous blocks to preserve within-block dependence; we use a block length of 168 hours (one week) and 300 bootstrap replications, and summarize uncertainty using the dispersion/quantiles of the bootstrap slope estimates [25][26].

5.4 Implementation details and reproducibility

All analysis was implemented in Python using XGBoost for nuisance estimation and scikit-learn utilities for tuning and cross-fitting. Randomness was controlled by fixing the random seed in all model training and resampling steps, and the full code used to generate the tables and figures is provided as supplementary material.

A key identification requirement is the overlap (positivity) assumption, meaning air density must retain meaningful variation after conditioning on the confounders.

__Table 4:__ Overlap (positivity) diagnostics: Std_RT range (standard deviation of treatment residuals, i.e., air density after removing variation explained by X) across wind-speed bins for both islands under scenarios A–D.

Scenario	Scenario label	North Island Std_RT range (kg/m³)	South Island Std_RT range (kg/m³)
A	Active-generation (base filtered)	0.0089–0.0126	0.0118–0.0131
B	Active-wind incl. zeros	0.0089–0.0128	0.0118–0.0131
C	All-hours (no wind filter)	0.0089–0.0140	0.0119–0.0155
D	All-hours excl. operational zeros	0.0089–0.0140	0.0120–0.0155

We check this using the standard deviation of the treatment-residual series (Std_RT) within wind-speed bins, reported in **Table 4** (overlap diagnostics). Across chosen sites from both islands and across wind regimes, the residual treatment variation remains non-zero, which supports the overlap condition needed for causal estimation after adjustment. This type of diagnostic aligns with standard practice for checking support/overlap issues in observational adjustment settings.

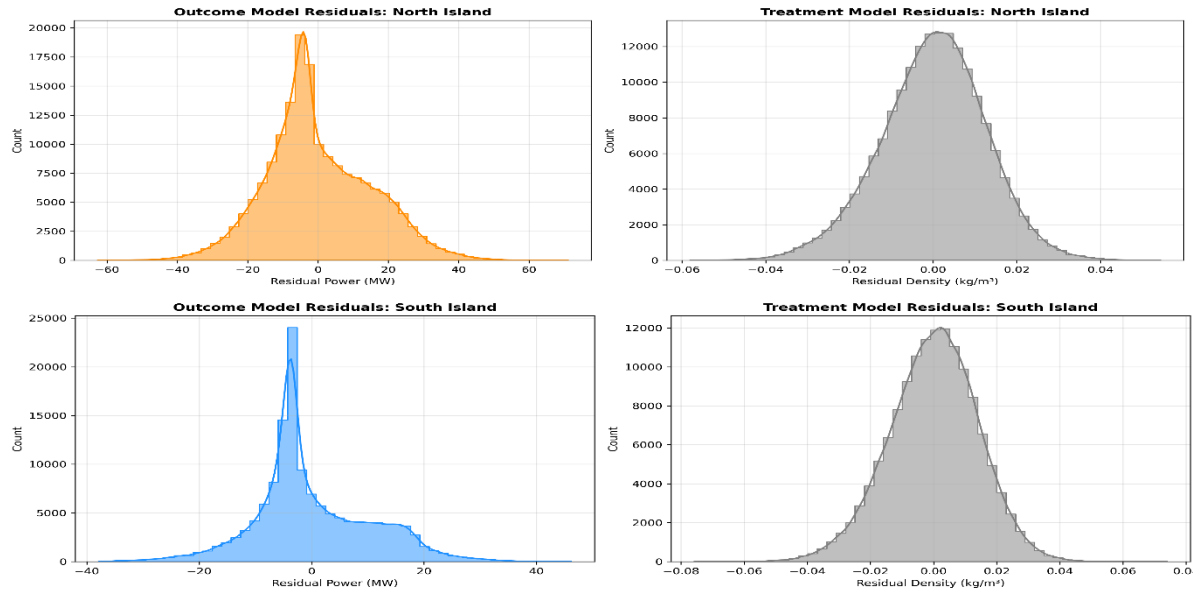

***Figure 10:** Residual diagnostics for the DML nuisance models: histograms of outcome residuals (power, left) and treatment residuals (air density, right) for Te Apiti wind farm (North Island) and White Hill wind farm (South Island)*

Finally, we assess nuisance-model adequacy by inspecting the residual distributions for both the outcome and treatment models (**Figure 10**). Treatment residuals are tightly centered around zero, while outcome residuals are broader with heavy tails, which is expected given unmodeled operational noise and rare events in generation. These diagnostics indicate that the first-stage models capture the dominant structure, leaving residual variation suitable for the final residual-on-residual regression step.

6. Results

This section reports the results from the Double Machine Learning analysis. We first present the overall causal estimates for the Te Apiti wind farm (North Island) and White Hill wind farm (South Island), then examine the temporal structural break in the White Hill wind farm (South Island). Finally, we report heterogeneity across seasonal and operating subgroups and use SHAP to summarise which inputs drive the fitted nuisance models.

6.1 Main Causal Estimates

Table 5: *Compact summary of DML results for each site: overall air-density effect with E-value, plus seasonal/diurnal/wind-regime ranges and the peak subgroup effect.*

Region	Overall effect (MW per 0.1 kg/m3)	E-value	Seasonal range (min–max)	Diurnal (Night–Day)	Wind regime range (min–max)	Peak subgroup (effect)
Te Apiti wind farm (North Island)	17.21 [16.63, 17.78]	29.98	10.25–26.48	15.87–18.54	10.53–30.43	Summer (26.48); Medium wind 6–10 m/s (30.43)
White Hill wind farm (South Island)	3.88 [3.49, 4.27]	10.00	0.21-5.91	2.40–5.08	-0.42 – 13.1	Autumn (5.91); Medium wind 6–10 m/s (13.1)

The main aim of this study was to estimate the marginal causal effect of air density on wind power after adjusting for meteorological confounding. Using the DML setup with time-ordered cross-fitting, we first partial out wind speed and seasonal structure through nuisance models and then estimate the density effect from the residual-on-residual regression. The results are summarised in **Table 5** and show a clear positive causal effect for both islands, consistent with the expected physics direction.

For the North Island (Te Apiti), the estimated coefficient is 17.21 MW per 0.1 kg/m3 (95% CI: 16.63 to 17.78; $p < 0.001$). A 0.1 kg/m3 change in air density therefore corresponds to about a 17.2 MW change in expected generation, which is meaningful against a site capacity of about

90 MW. The corresponding E-value is 29.98, suggesting that only very strong unmeasured confounding would be able to fully explain away the observed effect.

For the South Island (White Hill), the estimated effect is smaller at 3.88 MW per 0.1 kg/m3 (95% CI: 3.49 to 4.27; $p < 0.001$), with an E-value of 10.00. This indicates that air density remains a statistically robust driver in both sites, but the marginal sensitivity is substantially lower in the South Island sample. The cross-fitting step reduces overfitting by estimating nuisance predictions on held-out folds before the final regression, following standard DML practice [27].

6.1.1 White Hill wind farm (South Island): The Temporal Structural Break

Table 6: *White Hill wind farm (South Island) era diagnostics (Early 2005–2014 vs Late 2015–2024): sample size and summary statistics for power, air density, wind speed, and seasonal composition.*

Era	N	Power mean (MW)	Power range (MW)	Density mean (kg/m³)	Density std	Wind Speed mean (m/s)	Summer (%)
Early (2005–2014)	51393	20.72	0.0–55.4	1.1939	0.0247	5.13	0.235
Late (2015–2024)	86702	16.93	0.0–54.9	1.1895	0.0249	5.97	0.248
All (2005–2024)	138095	18.34	0.0–55.4	1.1911	0.0249	5.66	0.243

The White Hill wind farm (South Island) analysis shows a clear temporal break when the data are split into an Early era (2005–2014) and a Late era (2015–2024). As reported in **Table 6**, the Early era has fewer valid observations (N = 51,393) than the Late era (N = 86,702), even though both cover roughly a decade. This indicates a change in data availability and operating conditions across time, so the pooled estimate should be interpreted as a long-run average.

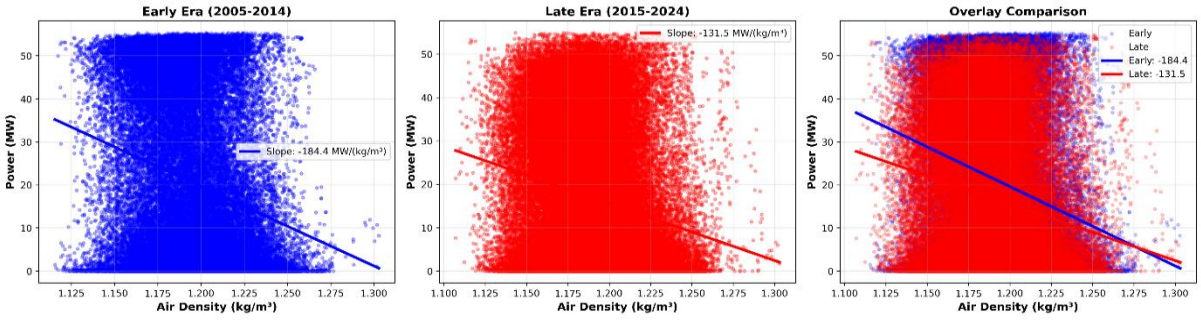

Figure 11: *White Hill wind farm (South Island) density–power relationship by era (Early 2005–2014 vs Late 2015–2024), showing the unadjusted fitted slope in each period and an overlay comparison.*

Figure 11 highlights the shift in the density–power relationship: the fitted unadjusted slope is more negative in the Early era (−184.4 MW/(kg/m³)) than in the Late era (−131.5 MW/(kg/m³)).

Table 7: White Hill wind farm (South Island) density–power correlations by era, showing that the raw correlation is negative but the wind-speed-adjusted (partial) correlation flips positive in the Late era.

Era	Raw corr	Raw p	Partial corr (controls: wind speed)	Partial p	Sign flip
Early (2005–2014)	-0.2198	0	0.0064	0.1495	YES
Late (2015–2024)	-0.2089	0	0.0806	0	YES

Table 7 supports the same confounding story: raw correlations remain negative in both eras, but the partial correlation (controlling for wind speed) is near zero and not significant in the Early era ($r = 0.0064$, $p = 0.1495$) and becomes positive and significant in the Late era ($r = 0.0806$, $p < 0.001$). This pattern is consistent with covariate shift over time, where the joint distribution of inputs and operating context changes, affecting simple associations.

6.2 Heterogeneity Analyses

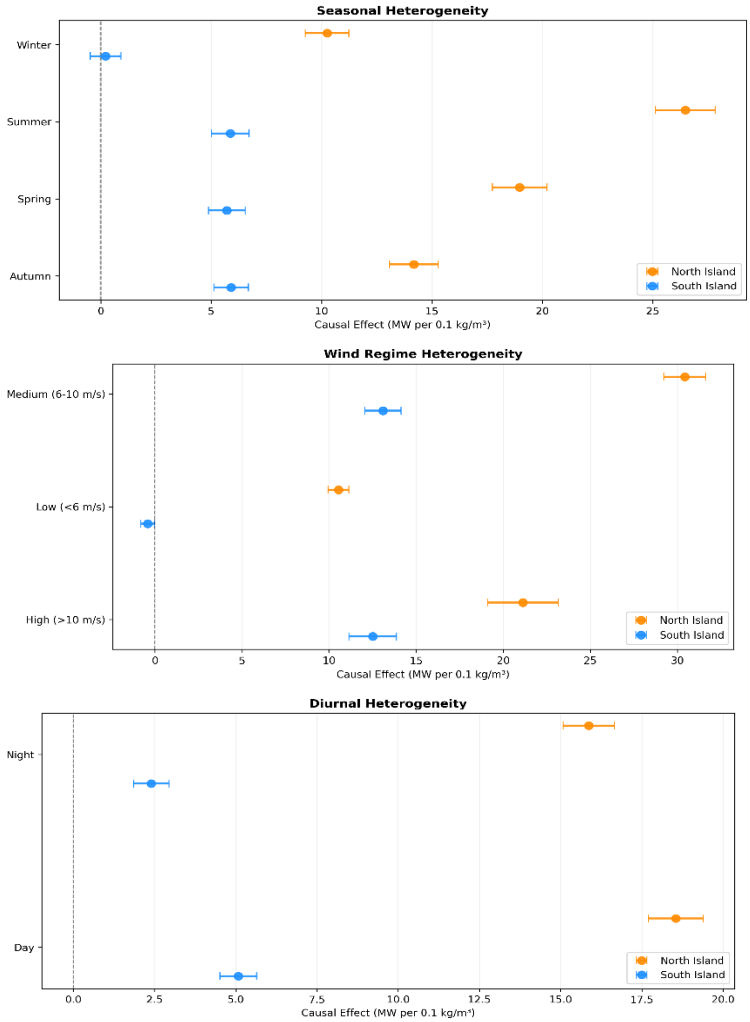

***Figure 12**. Heterogeneous Treatment Effects (HTE) of air density on wind power generation across three operational dimensions. Estimates represent the causal effect in MW per 0.1 kg/m3 with 95% confidence intervals computed using the*

normal approximation ($\pm 1.96 \times SE$) from the residual-on-residual regression. (Top) Seasonal heterogeneity showing the "marginal gain" effect in summer versus saturation in winter. (Middle) Wind regime heterogeneity showing peak sensitivity in the "ramp" region (6–10 m/s). (Bottom) Diurnal heterogeneity showing a consistent daytime premium due to vertical mixing. Orange circles denote North Island (Te Apiti); Blue circles denote South Island (White Hill).

To understand the operational conditions under which air density most strongly influences generation, we estimated Heterogeneous Treatment Effects (HTE) across season, wind regime, and time of day; subgroup coefficients with 95% confidence intervals are summarized in **Figure 12.**

Seasonal analysis indicates a clear “marginal gain” pattern. In the Te Apiti wind farm (North Island), the causal effect peaks in summer at ~26.48 MW per 0.1 kg/m³ and is lowest in winter at ~10.25 MW per 0.1 kg/m³, consistent with density being most valuable when conditions are closer to the production margin rather than saturation. In the South Island, the winter estimate is small and statistically indistinguishable from zero (0.21 MW per 0.1 kg/m³; 95% CI spans negative to positive), consistent with high-variability winter flow regimes where operational controls (e.g., storm-related cut-out/curtailment) can dominate realized output. This type of subgroup variation is consistent with treatment heterogeneity issues commonly encountered in complex systems [28].

Stratification by wind speed supports a mechanically interpretable regime dependence. Both sites show peak sensitivity in the medium wind band (6–10 m/s): ~30.43 MW per 0.1 kg/m³ (Te Apiti wind farm in the North Island) and ~13.10 MW per 0.1 kg/m³ (White Hill wind farm in the South Island). At high winds (>10 m/s), the effect dampens (~21.14 Te Apiti wind farm in the North Island; ~12.50 White Hill wind farm in the South Island per 0.1 kg/m³), consistent with saturation/pitch-regulation near rated operation, while the South Island low-wind regime (<6 m/s) shows a small negative estimate (−0.42 MW per 0.1 kg/m³), This negative effect may reflect control thresholds or rare-event behaviour under weak-wind, high-density conditions where the linear DML model may not capture the true conditional expectation. Alternatively, such regimes are infrequent in the South Island, and the estimate carries higher uncertainty (95% CI spans negative → positive). [29]

Finally, diurnal decomposition shows a consistent daytime premium: in the Te Apiti wind farm (North Island), day is ~18.54 vs night ~15.87 MW per 0.1 kg/m³ (~17% higher), and in the White Hill wind farm (South Island) day is ~5.08 vs night ~2.40 MW per 0.1 kg/m³ (~2.1×). This pattern plausibly reflects daytime boundary-layer mixing versus nighttime stability, which can modulate the effective air-mass structure interacting with the rotor [30].

6.3 Mechanism Analysis (SHAP)

To check whether the machine learning model learned physically sensible relationships rather than only fitting correlations, SHAP (SHapley Additive exPlanations) values were used to interpret the fitted model [31].

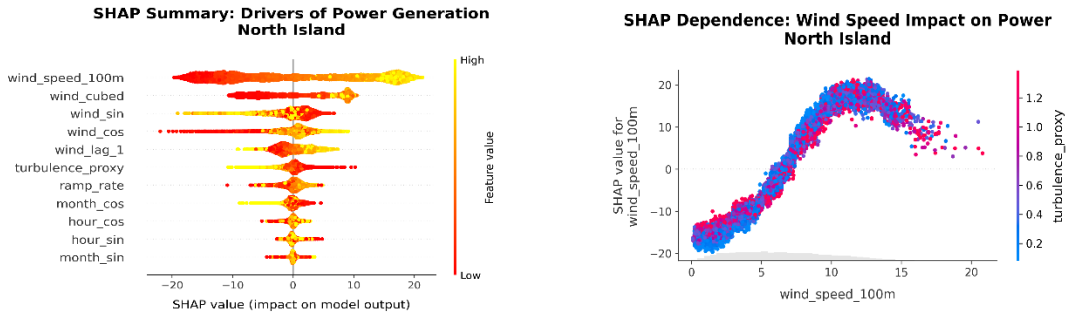

Figure 13. *SHAP interpretation for the North Island outcome model: (A) global feature importance (summary plot) and (B) wind-speed dependence showing a non-linear response and interaction with turbulence_proxy.*

The global summary plot (**Figure 13A**) shows that wind_speed_100m and wind_cubed are the strongest drivers of predicted power, which is consistent with wind power being dominated by wind availability. Other terms such as wind direction harmonics (wind_sin, wind_cos), lagged wind (wind_lag_1), and turbulence_proxy have smaller but consistent contributions, suggesting the model uses additional structure beyond instantaneous speed.

The dependence plot (**Figure 13B**) provides a shape check for wind speed. The SHAP contribution increases sharply from low winds into the mid-range and then flattens at higher wind speeds, which matches the expected ramp-to-saturation behaviour seen in empirical power curves. The colour overlay indicates that, at a fixed wind speed, higher turbulence_proxy values are often associated with a lower wind_speed_100m contribution, which is consistent with the idea that “wind quality” and short-term variability affect realised output [32].

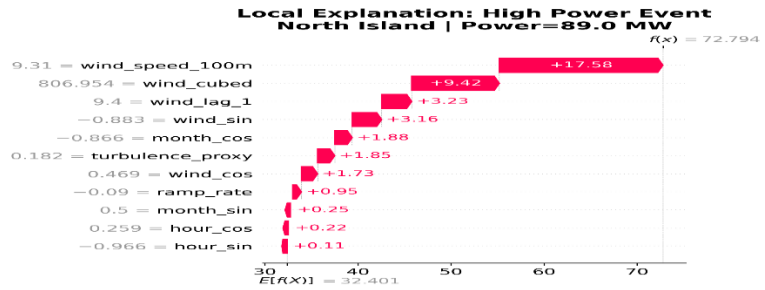

Figure 13C*. SHAP waterfall explanation for a representative high-power hour in the Te Apiti wind farm (North Island) outcome model, showing how individual features (e.g., windspeed100m, windcubed, windlag1, and directional terms) push the prediction from the baseline to the final output*

Finally, the waterfall plot (**Figure 13C**) illustrates the drivers behind a high-power hour (89.0 MW). Starting from a baseline expectation of 32.4 MW, the largest positive pushes come from wind_speed_100m (+17.58) and wind_cubed (+9.42), with additional support from wind_lag_1 (+3.23) and the directional terms. This event-level view makes it clear that the model's high-output predictions are mainly explained by strong wind conditions, with secondary adjustments from persistence and direction rather than unrelated variables.

7. Robustness and Validation

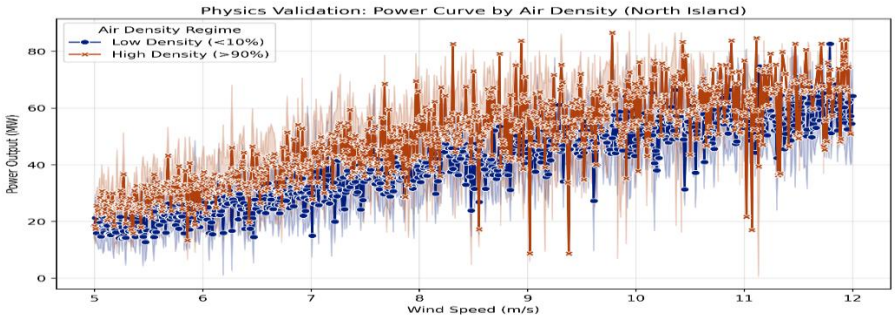

Figure 14: *Physics validation showing the empirical power curve stratified by air density. The vertical separation between high (dark) and low (light) density curves in the ramp region confirms the physical basis of the causal estimate.*

To ensure the estimated causal effects are not artifacts of model choice, preprocessing, or residual confounding, we conducted a set of validation and sensitivity checks.

7.1 Physics-Based and Statistical Validation

First, we performed a model-agnostic physics check using stratified empirical power curves. **Figure 14** compares the Te Apiti wind farm (North Island) power curve under low-density (10^{th} percentile) and high-density (90^{th} percentile) air density conditions. In the ramp region (about 5–12 m/s), the high-density curve generally sits above the low-density curve at similar wind speeds, which is consistent with the expected density scaling in aerodynamic power conversion [4].

Table 8: *Compact robustness summary for the air-density causal effect, showing placebo validity, late-era temporal stability (2015–2024), and learner sensitivity (Random Forest) for North and South Island models*

Check	North Island: Effect (p-value)	South Island: Effect (p-value)
Placebo (random noise treatment)	-0.012 (0.745)	-0.00008 (0.998)
Temporal stability (Late era, 2015–2024)	114.74 (<0.001)	56.57 (<0.001)
Learner robustness (Random Forest)	162.13 (<0.001)	34.74 (<0.001)

Note: Effects in Table 8 are raw DML coefficients (per 1 kg/m³) and are not scaled to per 0.1 kg/m³ as in Table 5. The purpose is to confirm sign consistency and statistical significance across specifications.

We tested estimator validity using placebo and robustness checks (**Table 8**). In the placebo test, replacing air density with random noise yields effects essentially equal to zero for both islands (Te Apiti wind farm in North Island: −0.012; White Hill wind farm in South Island: −0.00008) with clearly non-significant p-values, consistent with no spurious signal. The late-era stability check (2015–2024) retains a large positive estimated effect in both regions (Te Apiti wind farm in North Island: 114.74; White Hill wind farm in South Island: 56.57), with p-values reported <0.001, indicating strong statistical support. Learner robustness using Random Forest also preserves positive effects (Te Apiti wind farm in North Island: 162.13; White Hill wind farm in South Island: 34.74), suggesting the main finding is not driven by this particular modelling choice.

7.2 Sensitivity to Data Filtering and Unobserved Confounding

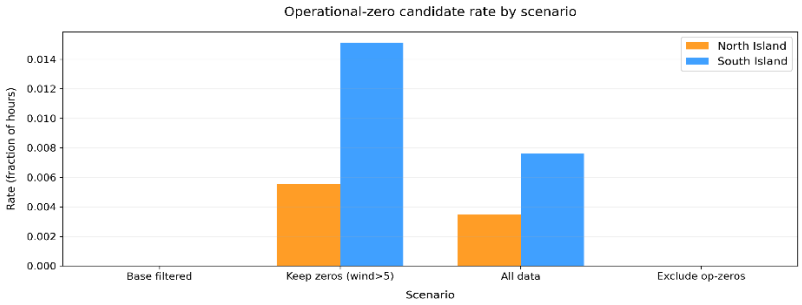

Figure 15: *Operational-zero candidate rate (fraction of hours) across preprocessing scenarios for the Te Apiti wind farm (North Island) and the White Hill wind farm (South Island).*

To test sensitivity to curtailment/operational-zero definitions, we compared multiple filtering scenarios. **Figure 15** reports the operational-zero candidate rate (fraction of hours) for Te Apiti wind farm (North Island) and White Hill wind farm (South Island) under Base filtered, keep zeros (wind > 5), All data, and Exclude op-zeros; rates rise when zeros are retained and drop to ~0 when operational zeros are excluded. This scenario contrast documents how treatment support changes under alternative curtailment filters and is used to motivate the positivity/overlap requirement in causal inference [12].

Potential unmeasured confounding was assessed using established causal inference principles [33]. The North Island E-value for the overall estimate is 29.98, implying that only very strong unmeasured confounding could explain away the effect, to be interpreted alongside double-robust estimation methods and related sensitivity guidance [34][35][36][37].

8. Discussion

This study shows that air density is a meaningful and physically interpretable driver of wind power generation in New Zealand, and it demonstrates why correlation alone can be misleading. Using a Double Machine Learning (DML) framework, the analysis corrects the negative density–power correlation seen in raw data—which occurs when high-pressure (high-density) regimes coincide with calmer winds—and instead estimates a robust positive causal effect after adjustment. The estimated overall effect is approximately 17.21 MW per 0.1 kg/m³ at the Te Apiti wind farm (North Island) (95% CI: 16.63–17.78; E-value: 29.98) and 3.88 MW per 0.1 kg/m³ at the White Hill wind farm (South Island) (95% CI: 3.49–4.27; E-value: 10.00). For operators such as Transpower, this suggests that density forecasts can add operational value beyond wind-speed-only baselines, consistent with system-level priorities around wind integration and forecasting quality.

8.1 Interpretation of the Conditional Coefficient

These estimates are conditional on the adjustment set (wind speed, seasonality, lagged terms, and harmonic features) and represent the residual density effect after removing confounding by meteorological structure through nuisance models (first-stage diagnostics: $R^2 \approx 0.64$). In this weather-adjusted setting, sensitivity is largest in the ramp region: heterogeneous effects peak at 30.43 MW per 0.1 kg/m³ for medium wind speeds (6–10 m/s), with an overall Te Apiti wind farm (North Island) wind-regime range of 10.53–30.43 MW per 0.1 kg/m³, explaining why the pooled effect remains operationally substantial.

8.2 Regional and seasonal heterogeneity

A second result is that Te Apiti (North Island) is substantially more sensitive to density than White Hill (South Island) (roughly 4–5× in the overall estimates). One plausible explanation is regional heterogeneity in fleet characteristics and control strategies—e.g., older stall-regulated or differently tuned turbines may transmit density variability into power more directly than modern pitch-regulated machines—yet this mechanism attribution should be treated as a working hypothesis unless supported by turbine-level metadata and site-specific operational information. Seasonal patterns reinforce a “marginal gain” interpretation: at Te Apiti, the seasonal effect ranges from 10.25–26.48 MW per 0.1 kg/m³, and at White Hill from 0.21–5.91 MW per 0.1 kg/m³, indicating that density is most valuable when turbines operate away from saturation.

8.3 Limitations and future directions

Firstly, synoptic state is not directly observed and is proxied through time features and wind speed, so unmeasured weather structure (e.g., boundary-layer stability, wind shear, fine-scale variability) may cause residual confounding; however, E-values (29.98 for Te Apiti; 10.00 for White Hill) suggest resilience to moderate unmeasured confounding. Secondly, the South Island exhibits temporal instability: the early era (2005–2014, N = 51,393) differs from the late era (2015–2024, N = 86,702) in mean power (20.72 vs 16.93 MW) and density distributions, so the pooled White Hill estimate (3.88 MW per 0.1 kg/m³) should be interpreted as a long-run average rather than a constant, while Te Apiti appears more temporally stable for operational use. Furthermore, results are based on representative single-site data (Te Apiti; White Hill) rather than fleet-wide generation, and the linear DML specification assumes additivity outside the reported subgroups; future work should incorporate turbine metadata, direct stability/shear proxies, and more flexible nonparametric causal estimators to test mechanisms and generalizability more rigorously, consistent with evidence that density variations matter for wind energy potential and should be treated explicitly in physics-aware analyses.

9. Conclusion

This study addresses a common issue in wind-power modelling: simple correlations between weather variables and generation can be confounded by synoptic regimes and seasonality. Using a Double Machine Learning (DML) setup on New Zealand grid wind generation data (Te Apiti wind farm (North Island) 2005–2024; South Island: 2009–2024) and reanalysis covariates, the causal effect of air density was estimated after adjusting for wind speed and seasonal structure, resolving the negative density–power association seen in the raw data.

Overall, a 0.1 kg/m³ increase in air density corresponds to about 17.2 MW higher expected generation in the Te Apiti wind farm (North Island) and about 3.9 MW in the White Hill wind farm (South Island). This regional gap is consistent with heterogeneity in fleet mix and operating regimes, but any attribution to specific turbine control technology (stall vs pitch) should be treated as a plausible explanation rather than a confirmed mechanism. The South Island also shows instability across an early vs late split, which points to changes in operating context and data coverage over time and reinforces the need for era-aware interpretation.

Operationally, the density effect is strongest in summer and at medium wind speeds (6–10 m/s), i.e., conditions closer to the ramp region. In higher-wind winter conditions, the marginal contribution of density is smaller because turbines more often operate near rated limits. Time-ordered cross-fitting and moving-block bootstrapping were used to respect temporal dependence in weather-driven data, and the moving-block bootstrap results were consistent with the analytical confidence intervals, indicating that serial dependence in the data does

not materially inflate the reported uncertainty. Placebo and learner checks further indicate that the estimated effect direction is stable under these validations.

CODE AVAILABILITY

The complete code and data processing framework for this study are available at: GitHub repo

REFERENCES

1 Electricity Authority of New Zealand, "Machine learning methods for predicting wind generation," Electricity Authority, Wellington, New Zealand, Tech. Rep., Oct. 2022. [Online]. Available: https://www.ea.govt.nz/documents/2385/Machine-learning-methods-for-predicting-wind-generation_MkxN3ZL.pdf

2 J. Pearl, "Causal diagrams for empirical research," Biometrika, vol. 82, no. 4, pp. 669–688, Dec. 1995, doi: 10.1093/biomet/82.4.669.

3 P. W. G. Tennant et al., "Use of directed acyclic graphs (DAGs) to identify confounders in applied health research: review and recommendations," Int. J. Epidemiol., vol. 50, no. 2, pp. 620–632, Apr. 2021, doi: 10.1093/ije/dyaa213.

4 A. Kalmikov, "Wind power fundamentals," in Wind Energy Engineering: A Handbook for Onshore and Offshore Wind Turbines, P. Veers, Ed. Amsterdam, Netherlands: Elsevier, 2017, pp. 17–24.

5 Lin, Chuan et al. "Causality analysis of meteorologically sensitive in offshore wind power generation based on the time-lagged convergent cross mapping." *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 238 (2024): 1299 - 1309.

6 Nyam Jargalsaikhan et al Exploring influence of air density deviation on power production of wind energy conversion system: Study on correction method, Renewable Energy, Volume 220, 2024, 119636, ISSN 0960-1481, https://www.sciencedirect.com/science/article/pii/S0960148123015513

7 Victor Chernozhukov et al Double/debiased machine learning for treatment and structural parameters, *The Econometrics Journal*, Volume 21, Issue 1, 1 February 2018, Pages C1–C68, https://doi.org/10.1111/ectj.12097

8 A. P. Sturman and N. J. Tapper, The Weather and Climate of Australia and New Zealand, 2nd ed. Melbourne, VIC, Australia: Oxford University Press, 2006.

9 Kidston, J., J. A. Renwick, and J. McGregor, 2009: Hemispheric-Scale Seasonality of the Southern Annular Mode and Impacts on the Climate of New Zealand. *J. Climate*, **22**, 4759–4770, https://doi.org/10.1175/2009JCLI2640.1.

10 Hersbach H, Bell B, Berrisford P, et al. The ERA5 global reanalysis. Q J R Meteorol Soc. 2020;146:1999–2049. https://doi.org/10.1002/qj.3803

11 Alain Ulazia et al Global estimations of wind energy potential considering seasonal air density changes, Energy, Volume 187, 2019, 115938, ISSN 0360-5442, https://doi.org/10.1016/j.energy.2019.115938

12 Zhu, Yaqian et al. “Core concepts in pharmacoepidemiology: Violations of the positivity assumption in the causal analysis of observational data: Consequences and statistical approaches.” *Pharmacoepidemiology and drug safety* vol. 30,11 (2021): 1471-1485. doi:10.1002/pds.5338

13 J. Pearl, M. Glymour, and N. P. Jewell, Causal Inference in Statistics: A Primer. Chichester, UK: Wiley, 2016.

14 R. Davidson and J. G. MacKinnon, Econometric Theory and Methods. New York, NY, USA: Oxford University Press, 2004.

15 J. D. Hamilton, Time Series Analysis. Princeton, NJ, USA: Princeton University Press, 1994.

16 Elwert, Felix. (2013). Graphical Causal Models. 10.1007/978-94-007-6094-3_13.

17 J. Pearl, Causality: Models, Reasoning, and Inference, 2nd ed. Cambridge, UK: Cambridge University Press, 2009.

18 Tyler J. VanderWeele, James M. Robins, Directed Acyclic Graphs, Sufficient Causes, and the Properties of Conditioning on a Common Effect, *American Journal of Epidemiology*, Volume 166, Issue 9, 1 November 2007, Pages 1096–1104, https://doi.org/10.1093/aje/kwm179

19 V. Chernozhukov, C. Hansen, N. Kallus, M. Spindler, and V. Syrgkanis, Applied Causal Inference Powered by ML and AI, 2024. [Online]. Available: https://causalml-book.org

20 Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16). Association for Computing Machinery, New York, NY, USA, 785–794. https://doi.org/10.1145/2939672.2939785

21 F. Pedregosa, et al "Scikit-learn: Machine learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

22 M. Optis and J. Perr-Sauer, "The importance of atmospheric turbulence and stability in machine-learning models of wind farm power production," Renew. Sustain. Energy Rev., vol. 112, pp. 27–41, Sep. 2019, doi: 10.1016/j.rser.2019.05.031.

23 P. Bach et al "Hyperparameter tuning for causal inference with double machine learning: A simulation study," Proc. Mach. Learn. Res., vol. 236, pp. 1065–1117, 2024. [Online]. Available: https://proceedings.mlr.press/v236/bach24a/bach24a.pdf

24 D. N. Politis and J. P. Romano, "The stationary bootstrap," J. Amer. Statist. Assoc., vol. 89, no. 428, pp. 1303–1313, Dec. 1994, doi: 10.1080/01621459.1994.10476870.

25 D. N. Politis and H. White, "Automatic block-length selection for the dependent bootstrap," Econometric Rev., vol. 23, no. 1, pp. 53–70, 2004, doi: 10.1081/ETC-120028836.

26 Hans R. Kunsch "The Jackknife and the Bootstrap for General Stationary Observations," The Annals of Statistics, Ann. Statist. 17(3), 1217-1241, (September, 1989)

27 P. Bach, V. Chernozhukov, M. S. Kurz, and M. Spindler, "DoubleML — An object-oriented implementation of double machine learning in Python," J. Mach. Learn. Res., vol. 23, no. 53, pp. 1–6, 2022. [Online]. Available: https://www.jmlr.org/papers/v23/21-0862.html

28 S. Wager and S. Athey, "Estimation and inference of heterogeneous treatment effects using random forests," J. Amer. Statist. Assoc., vol. 113, no. 523, pp. 1228–1242, 2018, doi: 10.1080/01621459.2017.1319839.

29 Susan Athey, Guido W. Imbens. 2019. Machine Learning Methods That Economists Should Know About. *Annual Review Economics*. 11:685-725. https://doi.org/10.1146/annurev-economics-080217-053433

30 Howland, M. F., Ghate, A. S., Quesada, J. B., Pena Martínez, J. J., Zhong, W., Larrañaga, F. P., Lele, S. K., and Dabiri, J. O.: Optimal closed-loop wake steering – Part 2: Diurnal cycle atmospheric boundary layer conditions, Wind Energ. Sci., 7, 345–365, https://doi.org/10.5194/wes-7-345-2022

31 S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. 31st Int. Conf. Neural Inf. Process. Syst. (NIPS), Long Beach, CA, USA, Dec. 2017, pp. 4765–4774.

32 C. Molnar, Interpretable Machine Learning: A Guide for Making Black Box Models Explainable, 3rd ed. Independently published, 2025. [Online]. Available: https://christophm.github.io/interpretable-ml-book/

33 M. A. Hernán and J. M. Robins, Causal Inference: What If. Boca Raton, FL, USA: Chapman & Hall/CRC, 2020.

34 H. Bang and J. M. Robins, "Doubly robust estimation in missing data and causal inference models," Biometrics, vol. 61, no. 4, pp. 962–972, Dec. 2005, doi: 10.1111/j.1541-0420.2005.00377.x.

35 Dorn, Jacob & Guo, Kevin & Kallus, Nathan. (2024). Doubly-Valid/Doubly-Sharp Sensitivity Analysis for Causal Inference with Unmeasured Confounding. Journal of the American Statistical Association. 120. 1-23. 10.1080/01621459.2024.2335588.

36 Tyler J. VanderWeele, Peng Ding. Sensitivity Analysis in Observational Research: Introducing the E-Value. Ann Intern Med.2017;167:268-274. [Epub 11 July 2017]. doi:10.7326/M16-2607

37 J. D. Y. Kang and J. L. Schafer, "Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data," Statist. Sci., vol. 22, no. 4, pp. 523–539, Nov. 2007, doi: 10.1214/07-STS227.